NASA TECHNICAL MEMORANDUM

NASA TM X-58038
January 1970

DEVELOPMENT OF GUIDANCE-MONITORING TECHNIQUES AND GUIDANCE-MONITORING EXPERIENCE DURING THE APOLLO 11 LUNAR DESCENT

Presented at
the AIAA Sixth Annual Meeting and Technical Display
Anaheim, California, October 20-24, 1969

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION
MANNED SPACECRAFT CENTER
HOUSTON, TEXAS

NASA TM X-58038

DEVELOPMENT OF GUIDANCE-MONITORING TECHNIQUES AND GUIDANCE-MONITORING EXPERIENCE DURING THE APOLLO 11 LUNAR DESCENT

Christopher C. Kraft, Jr., Charley B. Parker,
Howard G. deVezin, Jr., and Emil R. Schiesser
Manned Spacecraft Center
Houston, Texas

PRECEDING PAGE BLANK NOT FILMED.

CONTENTS

FIGURES

DEVELOPMENT OF GUIDANCE-MONITORING TECHNIQUES AND GUIDANCE-MONITORING EXPERIENCE DURING THE APOLLO 11 LUNAR DESCENT

By Christopher C. Kraft, Jr., Charley B. Parker,
Howard G. deVezin, Jr., and Emil R. Schiesser
Manned Spacecraft Center

INTRODUCTION

To determine the adequacy of onboard guidance and propulsion systems for completing critical maneuvers safely, some means of monitoring the onboard systems must exist in the event that it becomes necessary to switch to a backup guidance system or to initiate an abort. In the Apollo Program, the onboard and ground systems were used as an integrated system to achieve safe and successful missions. Because there are two independent onboard guidance systems on the Apollo lunar module, it was necessary to use the ground tracking capabilities as the third source for defining guidance acceptability. Techniques were developed for using successfully ground tracking for inflight real-time monitoring of critical spacecraft maneuvers. The guidance monitoring systems and techniques which were developed for the lunar-landing phase of the Apollo 11 mission are described in this paper.

LUNAR-LANDING OPERATIONAL SEQUENCE

Lunar-landing preparation begins approximately 8 hours before the desired landing time with the transfer of the crew to the lunar module (LM) to begin checkout and initialization. The period from undocking to landing (fig. 1) includes the command and service module (CSM) and LM undocking, CSM separation, descent orbit insertion (DOI), and powered descent maneuvers. Brief descriptions of these phases of the lunar mission are presented in the following sections.

Undocking and Separation

During the 13th orbit and 2.5 hours before landing, the LM and CSM are undocked by a physical unlatching of a spring-loaded mechanism that imparts a relative velocity of 0.7 ft/sec to the vehicles. Stationkeeping begins at a distance of 40 feet and the LM is rotated around the longitudinal axis for observation of the landing gear and LM exterior by the command module (CM) pilot. One-half hour after undocking (one

revolution before powered descent initiation (PDI)), the CSM pilot performs a maneuver of 2.5 ft/sec directed radially downward toward the center of the moon which increases the separation distance to 11 216 feet (1.8 nautical miles) at DOI.

Descent Orbit Insertion

The DOI maneuver occurs one-half revolution after separation (or one-half revolution before PDI) and places the LM on a Hohmann transfer orbit to carry it from a near-circular orbit of 60 nautical miles to an altitude to 50 000 feet. The DOI maneuver consists of a two-jet 7.5-second ullage burn followed by a 28.5-second retrograde burn of the descent propulsion system (DPS). During the first 15 seconds of the burn, the DPS throttle is set to 11.7 percent. After this trim period, the throttle is increased manually to 40 percent for the remainder of the burn. The maneuver is designed to locate the perilune to the east of the landing-site longitude by a central angle of approximately 14° for PDI.

Powered Descent

The powered descent maneuver, which is initiated near the perilune of the descent transfer orbit, consists of three operational phases (fig. 2): braking (for efficiency), approach (for crew visibility), and landing (for manual control to touchdown on the lunar surface). The transition from the braking to the approach phase is termed high gate, and the crossover from the approach to the landing phase is called low gate. The ullage manuever and the DPS engine gimbal-angle trim period occur at the constant inertial attitude required by the primary guidance and navigation system (PGNS) guidance equations at DPS ignition.

The braking phase is designed for efficient reduction of orbital velocity and therefore uses maximum thrust from the DPS for most of the phase; however, the DPS is throttled during the final 2 minutes for guidance control. The LM is in a windows-down attitude until an altitude of approximately 45 000 feet. The LM is then rotated around the vehicle thrust axis to a windows-up attitude to permit use of the landing radar beginning at an altitude of approximately 39 250 feet. The braking-phase guidance is based on quadratic acceleration equations. The guidance automatically switches from braking-phase targets to approach targets at a time to go (to achieve targets) of 60 seconds.

The approach phase begins at an altitude of approximately 7000 feet (high gate) and provides for visual monitoring of the approach to the lunar surface; that is, the guidance is targeted to provide spacecraft attitudes which permit crewmember visibility of the landing area through the forward window throughout this phase. The same quadratic acceleration guidance law used in the braking phase is used in the approach phase. Under automatic guidance, a vertical descent is initiated when the time to go reaches 10 seconds.

The landing phase or the vertical portion of the descent begins at an altitude of 150 feet (low gate) and terminates at touchdown on the lunar surface. Normally, a 3-ft/sec rate of descent is used throughout the vertical descent. The guidance is a nulling routine of the velocity error. As previously mentioned, the vertical descent

under automatic guidance commences at an altitude of 150 feet. Operationally, however, the approach phase is considered to be terminated at an altitude of 500 feet, at which point the pilot normally assumes manual control of the vehicle for final selection of the touchdown point and for landing.

The actual descent trajectory (corrected for initial down-range position error) for Apollo 11 from the start of the approach phase or 26 000 feet from the target is shown in figure 3. The latter portion of the approach phase and the landing phase are presented in more detail in figure 4. The point marked "P66 initiation" represents the point at which the pilot assumed manual control of the LM. The angles between the thrust axis of the LM, the local vertical (pitch), and the forward axis of the LM and the line of sight to the target as measured by the landing-point designator (LPD) are also indicated in figures 3 and 4.

During the approach phase, the onboard computer calculates the angle between the LM forward axis and the line of sight to the target and displays this information to the pilot. By referring to a grid scribed on the LM forward window, the pilot can identify the target point to which the computer is attempting to fly. If this target is unacceptable, the pilot can "redesignate" the target point by making inputs to the computer through the control handle. This capability allows the pilot to make sizable adjustments to the targeted landing point during the approach phase. After manual takeover, the pilot can further refine his landing point by manual control. Because of visibility restrictions, it is difficult to fly more than a few hundred feet short of the target after manual takeover. Conversely, it is easier to fly long as depicted in figure 4, which shows that the Apollo 11 pilot flew approximately 1000 feet down range from the initial target point. Late redesignations unfortunately result in fuel penalties which, for a given magnitude of target-point redesignation, increase as the landing phase progresses. For example, the 120 pounds of fuel normally budgeted for redesignation would provide a change in range of approximately 3000 feet at the beginning of the approach phase (25 000-foot range, 7000-foot altitude) compared to a change of less than 600 feet at a 2000-foot range and 550-foot altitude. On the Apollo 11 flight, 450 pounds of fuel were expended for manual range extension.

LUNAR MODULE SOURCES OF TRAJECTORY INFORMATION AVAILABLE DURING DESCENT

The PGNS on board the LM consists of a large-capacity digital computer and a three-gimbal stable platform which perform all the guidance and navigation functions necessary for a powered descent, ascent, and rendezvous. The backup or abort guidance system (AGS) consists of a smaller digital computer and a strap-down inertial platform. The AGS can perform an ascent and rendezvous but cannot control the powered descent. The PGNS is initialized by a combination of ground updating via the digital up link and onboard operations such as platform alinement. The backup guidance system is initialized normally from the PGNS but also can be initialized by manually loading the required data based on ground-supplied information. Current state (trajectory) information from both navigation systems is transmitted to the ground via the telemetry down link to provide the basis for trajectory monitoring during powered flight.

The landing-radar system is considered an integral part of the guidance and navigation system for controlling the lunar descent. The landing-radar system is a four-beam Doppler system (fig. 5); three beams measure velocity in three axes and the fourth beam measures the slant range to the surface. The landing-radar antenna has two operating positions to provide full coverage of the wide range of LM pitch attitudes during powered descent. The landing-radar information is sent to the PGNS during powered descent and is used to modify the PGNS current state estimate to conform to the measured landing-radar altitude and velocity. These computations are performed automatically after the pilot gives the PGNS an initial "proceed" instruction. Initial-state and lunar-terrain uncertainties and navigation dispersions make incorporation of landing-radar data mandatory for a successful landing.

The rendezvous radar is a continuous-wave coherent system which operates in the X-band frequency range and uses phase-lock and Doppler-shift techniques to measure range and range rate. Directly available outputs are range, range rate, azimuth, and elevation. The primary purpose of the rendezvous radar is to measure the relative trajectory information between the CSM and the LM during rendezvous.

DEVELOPMENT OF EARLY GUIDANCE-MONITORING SYSTEMS

Need for Completely Independent Trajectory Data Source

When only the onboard system is considered, two sources of trajectory information are available for trajectory monitoring — the PGNS and the backup guidance system or AGS. During powered flight, the guidance and navigation system in control of the vehicle (usually the PGNS) will always indicate a nominal or near-nominal trajectory because the system is both navigating and controlling the vehicle. As long as the noncontrolling navigation system agrees with the controlling system, it can be assumed that both are correct. However, as soon as one system disagrees with the other, an impasse results unless additional information is available to isolate the malfunctioning system. Thus it is necessary to provide a completely independent source of trajectory information to act as a "tie breaker" in the event a malfunction should occur in one of the onboard guidance systems.

Developmental Guidelines

Development of ground-based and onboard trajectory-monitoring techniques for use during the powered descent and ascent phases of the lunar-landing mission began in 1964. Initial studies were directed toward determining what the ground monitoring capabilities were, or could be, during these phases of the mission and how these capabilities could complement and supplement the LM onboard monitoring capability. The studies established two facts which were used in the detailed development work that followed.

1. The onboard monitoring must provide the capability to detect malfunctions that result in rapid trajectory deviations and/or immediate crew-safety problems. This monitoring must be available because of the inherent delays associated with

obtaining, processing, interpreting, and communicating the required data on the ground. Although the terms "rapid" and "immediate" were not quantitatively defined, it was generally accepted that any malfunction that would result in crew-safety problems within 30 seconds of onset could not be adequately handled by ground monitoring.

2. No current capability existed either on the ground or on board the spacecraft to detect insidious trajectory deviations that could result in an unsafe trajectory.

The first fact was readily acceptable because it encompassed the same philosophy that was used in all previous manned space flights. However, the second fact was discouraging. The studies were indicative that the Doppler or range-rate information as measured along the line of sight from the spacecraft to the ground receiving station would be of high quality, resulting in the possibility of extremely precise measurements, at least along that particular line of sight. The studies were indicative that noise and bias on the radar data (particularly angle data) at lunar distances were too great to support powered-flight monitoring and that very simple filtering methods were ineffective. Therefore, it was necessary to proceed with the development of ground monitoring techniques, assuming that independent ground-based trajectory information (except for range-rate data) would not be acceptable to support the powered-flight portions of the lunar-landing mission unless very complex tracking data filters were used.

Early Guidance-Monitoring Systems

An analysis was made of all available sources of trajectory-type data and three readily available sources were identified: (1) the ground-based measurement of the range-rate information along the line of sight between the spacecraft and the ground receiving station, (2) range-rate information along the line of sight between the LM and the CM as measured by the LM rendezvous radar, and (3) landing-radar altitude and velocity information measured relative to the lunar surface during powered descent.

EVALUATION OF EARLY GUIDANCE-MONITORING SYSTEMS

None of the three possible "tie breakers" were entirely satisfactory. The limitations and application to guidance monitoring of each are discussed in the following sections.

Ground-Based Measurements

As noted previously, the only available ground-based trajectory information of sufficient accuracy to act as a "voter" was the range rate measured along the earth (tracker)/LM line of sight. The Manned Space Flight Network (MSFN) range-rate processor developed for this purpose used two sequentially telemetered state vectors from both guidance and navigation systems to determine equivalent velocity vectors along the earth (tracker)/LM line of sight (fig. 6). These vectors were then differenced from the MSFN measured velocity along the same line of sight. The differences were plotted as a function of time shown in figure 7. In the case shown, MSFN confirmed that the PGNS was operating normally even though the backup guidance system (AGS) appeared to be malfunctioning. Although the range-rate processor is capable of detecting very small

errors in velocities along the earth line of sight, it is insensitive to malfunctions which affect only velocity vectors perpendicular to the earth/LM line of sight (fig. 6). This limitation is significant for descents to landing sites near zero longitude when the downrange velocity vector lies essentially in this plane.

Lunar Module Rendezvous-Radar Range-Rate Measurements

The use of the LM rendezvous-radar data to isolate guidance systems malfunctions is essentially the same as that described for the MSFN data; that is, the calculated range rate between the LM and the CM based on the state vectors from each guidance system is differenced with the actual measured range rate and the difference plotted as a function of time. This system is limited by the same insensitivity to errors perpendicular to the line of sight, and additionally is constrained by the fact that the CM must be within approximately 400 nautical miles of the LM for the rendezvous radar to lock on and track. Another disadvantage of this system is the considerable amount of onboard computer time required to perform the monitoring functions.

Lunar Module Landing-Radar Measurements

The use of the landing-radar data to isolate guidance systems malfunctions is also similar to those previously described. In this instance, the PGNS and AGS state estimates were converted to equivalent landing-radar measurements of altitude and velocity components relative to the lunar surface and compared with the values for these parameters as measured by the landing radar. This system is limited from several standpoints, however. First, if an abort to lunar orbit during the descent phase became necessary, the landing radar would not be available during the powered ascent and a changeover to a different guidance-monitoring system would be required. Secondly, the landing radar is not, in the true sense, an independent measurement source because the output is used to modify the PGNS state estimate. The possibility exists that the landing radar could develop a malfunction which would result in feeding incorrect data into the PGNS state estimate. The PGNS state estimate would unknowingly converge on an invalid landing-radar estimate of state. The erroneous conclusion would be that the PGNS and landing radar agreed and were thus correct when, in fact, both would be wrong.

OPERATION OF EARLY GUIDANCE-MONITORING SYSTEM

Because none of the available trajectory "voters" were entirely satisfactory, it was decided to implement all three sources with the intention to use the best characteristics of each source in the final monitoring techniques.

As the technique development proceeded, the method determined to be most suitable for monitoring differences between the navigation states of the two guidance and navigation systems was to take the navigation state vectors from the respective guidance and navigation systems, time synchronize them by linear interpolation, transform the velocity components into a common coordinate system, and then difference them. The difference was displayed on strip-chart recorders for real-time analysis. The

common coordinate system chosen for displaying the differences defined one component as being parallel to the local horizontal (V_x), one component as being along the local vertical (V_z), and one component as being $V_x \times V_z$ (V_y). The system, named the local-horizontal/local-vertical coordinate system, is illustrated in figure 8. A display format which was readily interpreted and which was adaptable to the use of fixed-limit lines was developed to present these velocity errors to the flight controllers to assist them in determining the go/no-go status of each guidance system. Thus, the defined guidance technique is to use ground monitoring for any significant velocity differences between the two guidance systems. As long as no differences are observed, both systems are considered to be functioning properly. In the event of a disagreement, one of the three validation or "voter" sources is selected to determine which guidance system is malfunctioning. If it can be determined that the malfunction is severe enough to impair the ability of the guidance system to perform assigned tasks satisfactorily, a no-go is declared.

APOLLO 11 DESCENT GUIDANCE-MONITORING SYSTEMS

Generally, it was considered that the techniques previously described were not completely satisfactory for the trajectory monitoring function, and investigation of additional techniques that could perform these functions more satisfactorily was continued.

Powered Flight Processor

One area receiving special emphasis concerned a mathematical technique which used MSFN S-band Doppler data to compute the position and velocity of the LM during powered descent. This technique differed in two ways from the other monitoring systems being implemented: the entire LM state was computed and only ground-based tracking information was used. It was, therefore, a completely independent monitoring system. This new MSFN processor, named the powered flight processor (PFP), is a complex Kalman filter with a 21-state variable solution. The accuracy of the PFP in computing an estimate of the LM position and velocity is influenced by the number and respective geometry of the MSFN tracking stations being used at a given time. Three stations with good north-south and east-west geographic separation are necessary for the program to sense accurately both the in-plane and out-of-plane motion of the LM. The program is designed to process data from four stations to make the PFP more stable and less sensitive to variations in the data quality.

The following is a brief description of the computational algorithm of the PFP.

1. The processor uses its current estimate of the LM position to compute the change in range between the vehicle and a MSFN tracking station over a specified time interval. These computations are performed for one to four tracking stations (nominally three or four).

2. The change in range from the vehicle to the tracking station is then converted to Doppler count by a conversion factor.

3. These values of computed Doppler count are differenced with the observed values of Doppler count measured by the tracking station over the same time interval. The differences are called residuals.

4. New values of the LM position and velocity and other state elements which minimize the residuals are computed.

5. The new values of LM position and velocity and the associated weighting matrix (a covariance matrix) are mapped forward to the time of the next set of Doppler observations, and the process is repeated.

This update-map cycle of the program is a continuous process done without manual interruption for the complete descent phase.

In addition to solving for LM position and velocity, the PFP also solves for LM thrust direction and rate of change of thrust direction, LM mass and mass flow rate, specific impulse, variable Doppler-rate biases, and integration constants used in the Doppler residual computations.

Final Apollo 11 Guidance-Monitoring Systems

The PFP and the MSFN range-rate processor discussed previously were implemented to support the trajectory-monitoring functions of the Apollo 11 lunar-landing mission. It should be noted that both of these systems were also used to perform the monitoring function for the ascent phase of the mission. The landing-radar data processor was used during Apollo 11 only to monitor the performance of the landing radar. The rendezvous-radar data processor was deleted altogether.

GUIDANCE MONITORING BY THE USE OF POWERED FLIGHT PROCESSOR

Pre-Apollo 11 Flight Experience

The Apollo 10 lunar-orbit rendezvous mission provided an almost perfect opportunity to "sled test" the PFP under actual lunar-environment conditions. After completion of the rendezvous exercise in the Apollo 10 mission, the ascent stage of the LM was jettisoned from the CM. The ascent-stage engine then was fired by ground command for a period of approximately 4 minutes. This burn pushed the ascent stage into a solar orbit and provided flight controllers the opportunity to operate the PFP and evaluate its performance during a long-duration powered-flight maneuver.

A subsequent analysis of the performance of the processor during this maneuver showed that the PFP was extremely effective and accurate, and a high degree of confidence in its capability as a guidance monitoring system was achieved.

Operational Performance During Apollo 11 Descent

During the Apollo 11 powered descent and landing, the PFP performed with consistent accuracy and stability as a real-time trajectory-monitoring system at lunar distances, as shown in figures 9 to 12. The parameters presented in these figures include LM altitude, altitude rate, pitch angle, and mass flow rate as functions of time of ignition (T_{ig}) for the powered descent. The circled data values for altitude and altitude rate during the 4 minutes of free flight before ignition (figs. 9 and 10) were determined predescent from the vector used to initialize the PFP and the onboard telemetry systems. The X's were plotted in real time.

The performance of the PFP was nominal until shortly after ignition when the MSFN data were lost for approximately 2 minutes. At the end of this interval, it was necessary to restart the filter with the PGNS estimate of the LM position. After the restart, the filter again performed nominally, which was indicative of close agreement with the premission nominal trajectory.

The velocity differences between the PFP and the PGNS as actually observed during the descent are presented in figure 13. Of particular interest is the $\Delta\dot{Z}$ or altitude-rate (vertical velocity) trace which indicates that the PFP measured a vertical descent velocity that was 19 to 20 ft/sec greater than that measured by the PGNS during the first 6 minutes of the descent burn. This error corresponds to the difference between the predescent and real-time values of altitude rate computed by the PFP (fig. 9) for the same period.

This condition was diagnosed correctly by the flight controllers as a down-range position error in the PGNS. (The same condition also was noted in the AGS, because the AGS was initialized from the PGNS.) As shown in figure 14, a central angle down-range position error θ manifests itself as an error in vertical velocity. There is very little difference between the magnitude of V and V', the state velocity of the PGNS and PFP systems, respectively. Therefore, there is very little difference between the $\dot{X}$ or down-range velocity components because for small angles of θ, $\cos\theta$ is very nearly 1. The $\dot{Z}$ velocity difference, on the other hand, is a function of $V \sin\theta$ because V' is nearly along V_x. Because V is large, the resulting error in vertical velocity is sizable. For small values of θ, the error or difference in vertical velocity is a direct function of the down-range position error. It can be shown that the functional relationship between $\Delta\dot{Z}$ and the down-range position is approximately 1 ft/sec equivalent to 1000 feet. Thus, the measured value of $\Delta\dot{Z}$ corresponded to an error in the down-range position of approximately 20 000 feet at the start of powered descent. As shown in figure 13, the vertical-velocity error decreased in magnitude to approximately zero between 6 and 8 minutes after the start of powered descent. This decrease was the result of landing-radar velocity measurements being incorporated into the PGNS state estimate which corrected the PGNS velocity estimates to agree with those of the PFP. The PFP was operated in a special way in order that the altitude rate error would be related to a down-track position error. The abort limit lines were then set accordingly.

Accuracy of Down-Range Estimates

It should be noted that the landing-radar update did not correct the PGNS estimate of down-range position but only its knowledge of the velocity and altitude of the LM relative to the lunar surface. The PGNS estimated landing point still was approximately 20 000 feet in error at landing. It should also be noted, however, that an error of this seemingly large magnitude is explained easily. In the nominal Apollo flight plan for lunar-descent preparations, the LM is tracked on one earthside pass of the moon and updated on the next earthside pass. Powered descent then is performed during the next pass. As a result, the last time that the PGNS state estimate normally is updated by the ground is approximately 2 hours (one orbit) before powered descent. In turn, this state estimate is based on tracking data that are also 2 hours old. Consequently, an error in state velocity as small as 0.5 ft/sec propagated over this 4-hour period results in more than 20 000 feet of down-range position error. Because landing-point accuracy was a secondary consideration on the Apollo 11 mission, this down-range position error was unimportant. However, for future missions which will have the objective of precision landing in or at a particular area or point, emphasis will be placed on eliminating the sources of error which degrade the PGNS state estimate. The PGNS-versus-PFP differences in down-range velocity ($\Delta\dot{X}$) and cross-range velocity ($\Delta\dot{Y}$) were well within the expected tolerances (fig. 13), confirming excellent operation of both the PFP and PGNS systems.

GUIDANCE MONITORING USING THE MSFN RANGE-RATE PROCESSOR

The MSFN range-rate processor also detected the down-range position error in the PGNS state. Apollo 11 descent range-rate residuals as shown in figure 15 indicate that this processor was measuring a difference of approximately 14 ft/sec along the MSFN/LM line of sight at the start of powered descent. This difference is consistent with the corresponding errors in PFP measurements shown in figures 9 and 13. As with the PFP data, the difference remained approximately constant until landing-radar updating, at which time the PGNS difference decreased to zero while the AGS (which is not updated by the landing radar) continued at about the same level. The geometry which allows the range-rate residuals to "see" the down-range position error is shown in figure 16. It should be noted that when the central angle $\emptyset$ is small, $\dot{R}_{MSFN} - \dot{R}_{onboard} \approx -V \sin \theta$ which can be recognized as the PFP $\Delta\dot{Z}$ equation. For Apollo 11, $\emptyset$ was approximately 35°, and the MSFN range-rate processor could "see" only approximately 80 percent of the effective Z velocity error.

CONCLUDING REMARKS

The PFP has proved to be an excellent processor for monitoring powered flight at lunar distances and will be the prime external source of trajectory information on future lunar-landing missions. To ensure the most effective use of the PFP, the powered-flight maneuvers should be planned for good geometrical coverage by at least three (and preferably four) MSFN sites during the maneuver. For future missions, it

appears possible that PFP data can be used to improve down-range position accuracy by updating the primary guidance system state estimates at or near the initiation of the descent phase.

The MSFN range-rate processor will be retained as a backup guidance-monitoring system for future missions. Monitoring of landing-radar data will continue but will be concerned with verifying the landing-radar performance rather than concerned with guidance monitoring. As a result of the development of the PFP, the rendezvous-radar range-rate processor will not be retained as a guidance-monitoring system because of the limited capability and excessive computer-loading requirements.

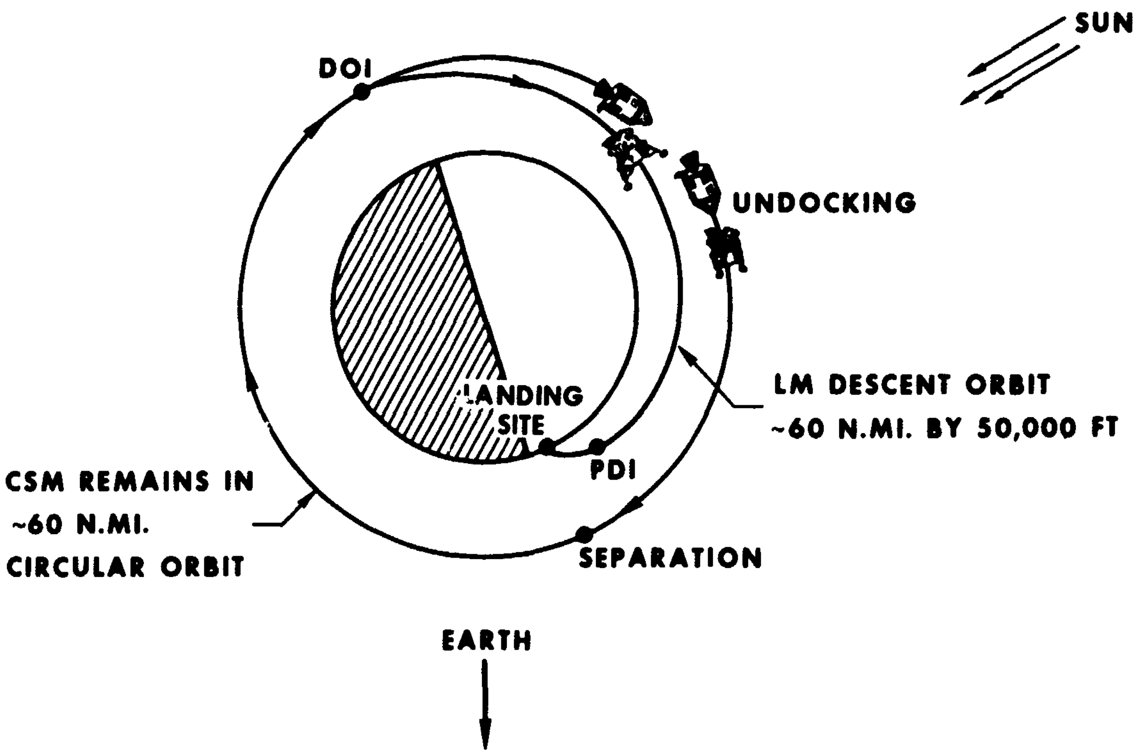

Figure 1. - Nominal descent profile.

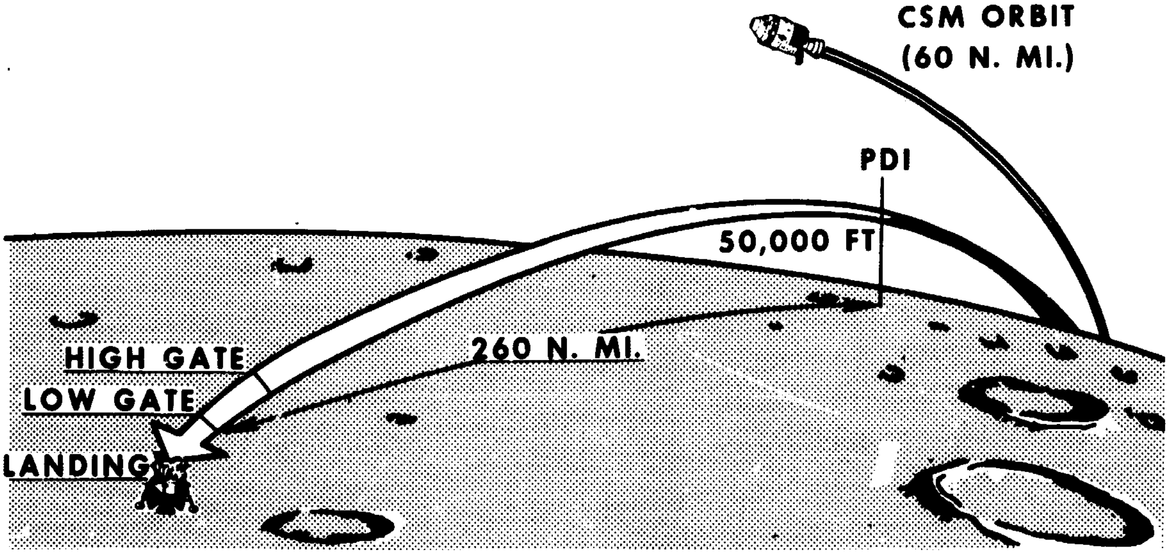

PHASE	INITIAL EVENT	DESIGN CRITERIA
BRAKING	PDI	MINIMAL PROPELLANT USAGE
APPROACH	HIGH GATE	CREW VISIBILITY
LANDING	LOW GATE	MANUAL CONTROL

Figure 2. - Operational phases of powered descent.

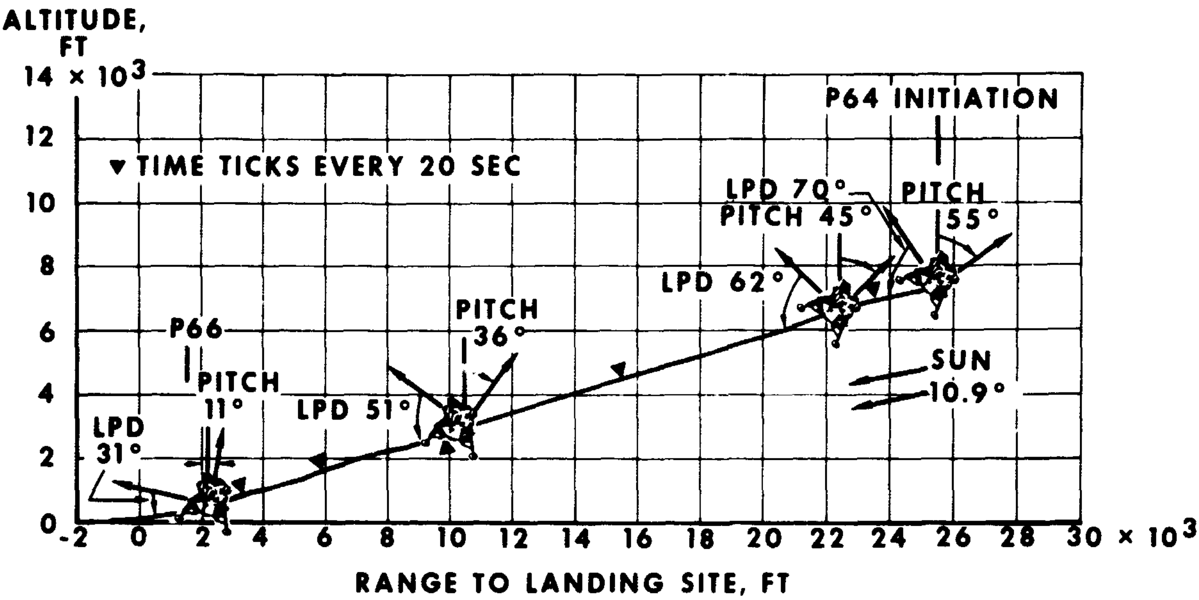

Figure 3. - Apollo 11 approach phase.

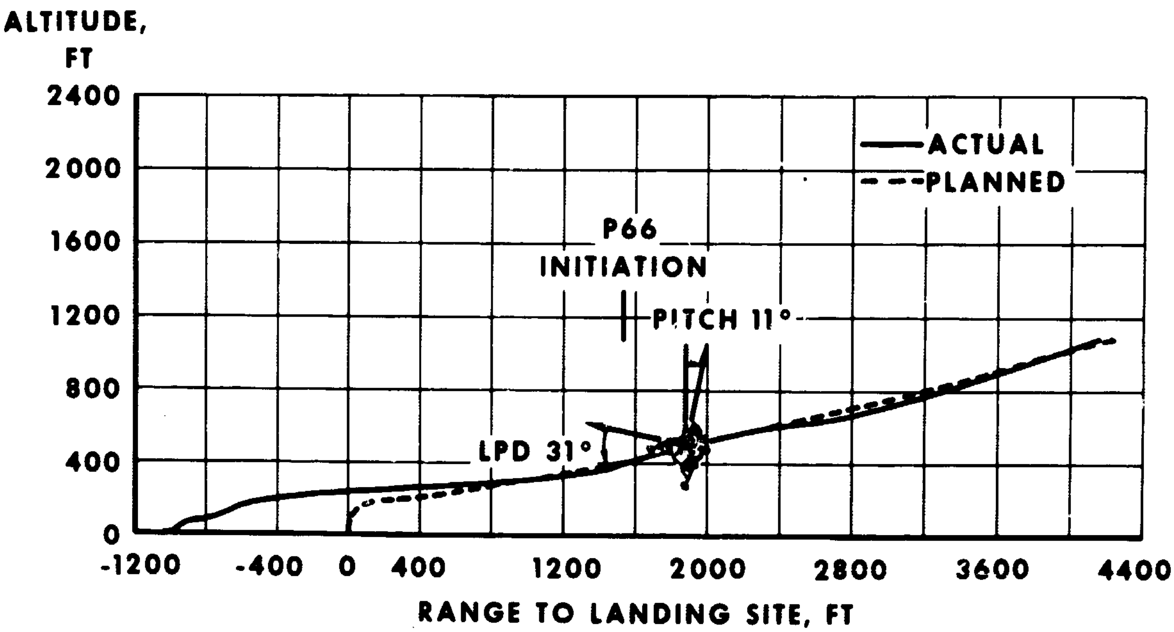

Figure 4. - Apollo 11 landing phase.

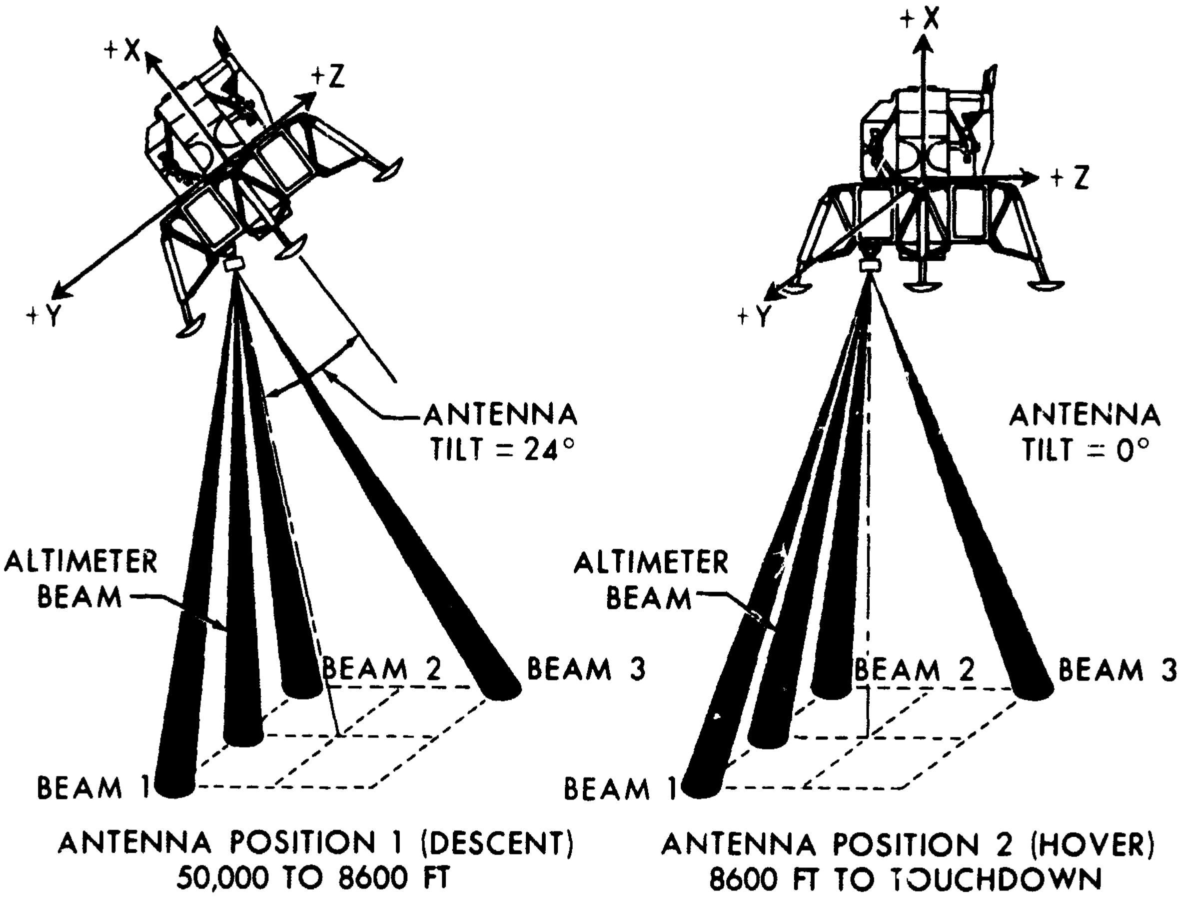

Figure 5. - Landing-radar beam configuration and antenna tilt angles.

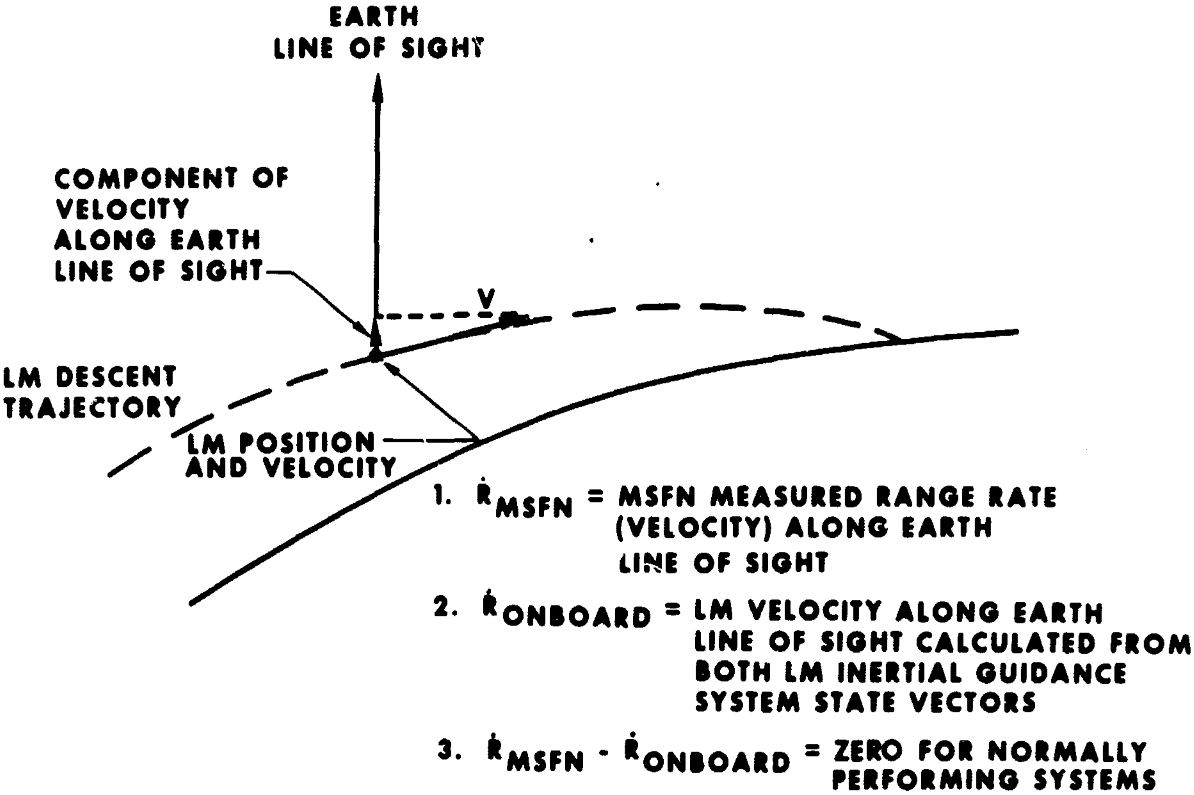

Figure 6. - Lunar module/earth geometry as viewed from the lunar north pole.

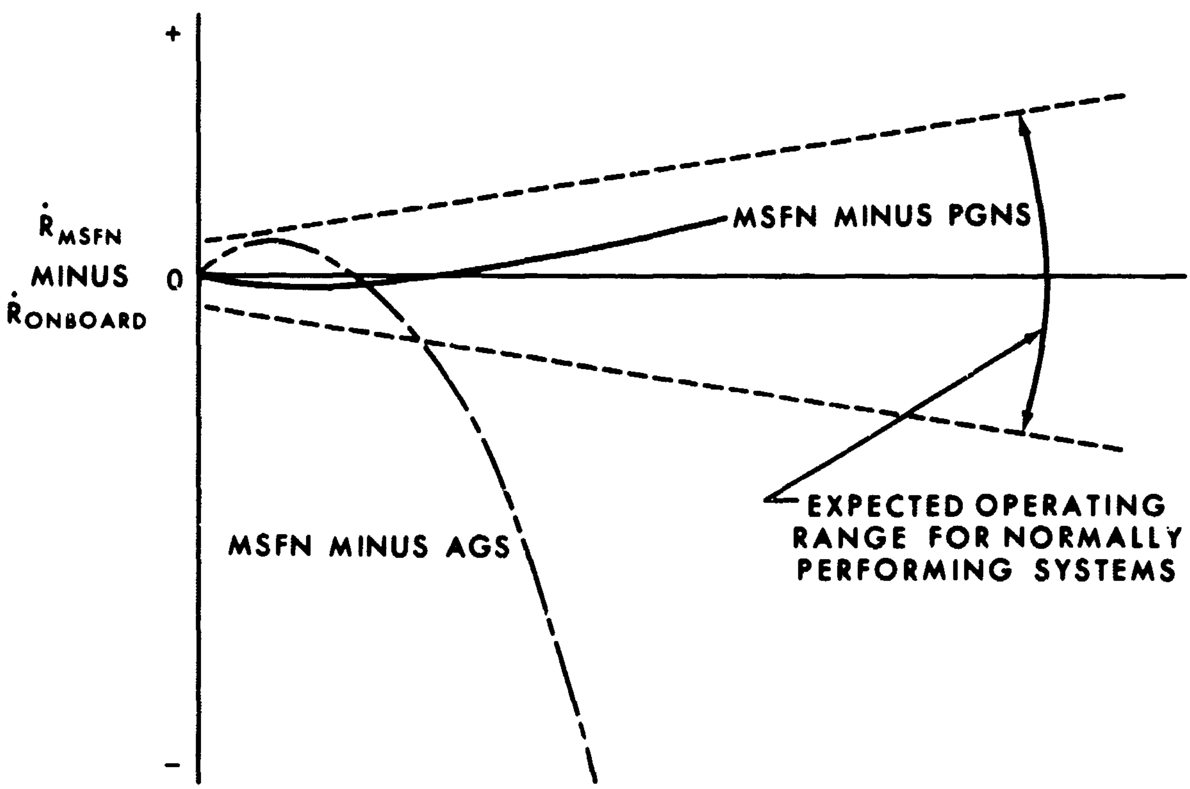

Figure 7. - Primary and abort guidance MSFN validation.

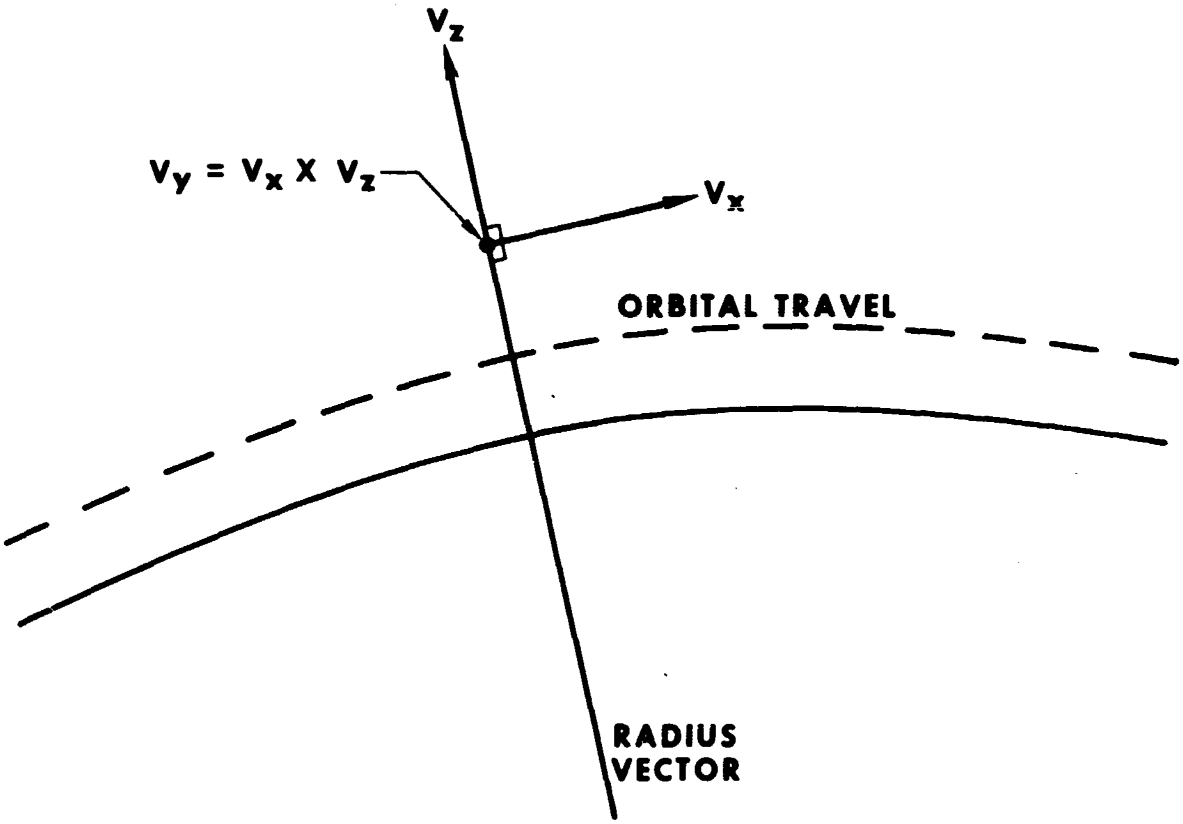

Figure 8. - Local-horizontal/local-vertical coordinate system.

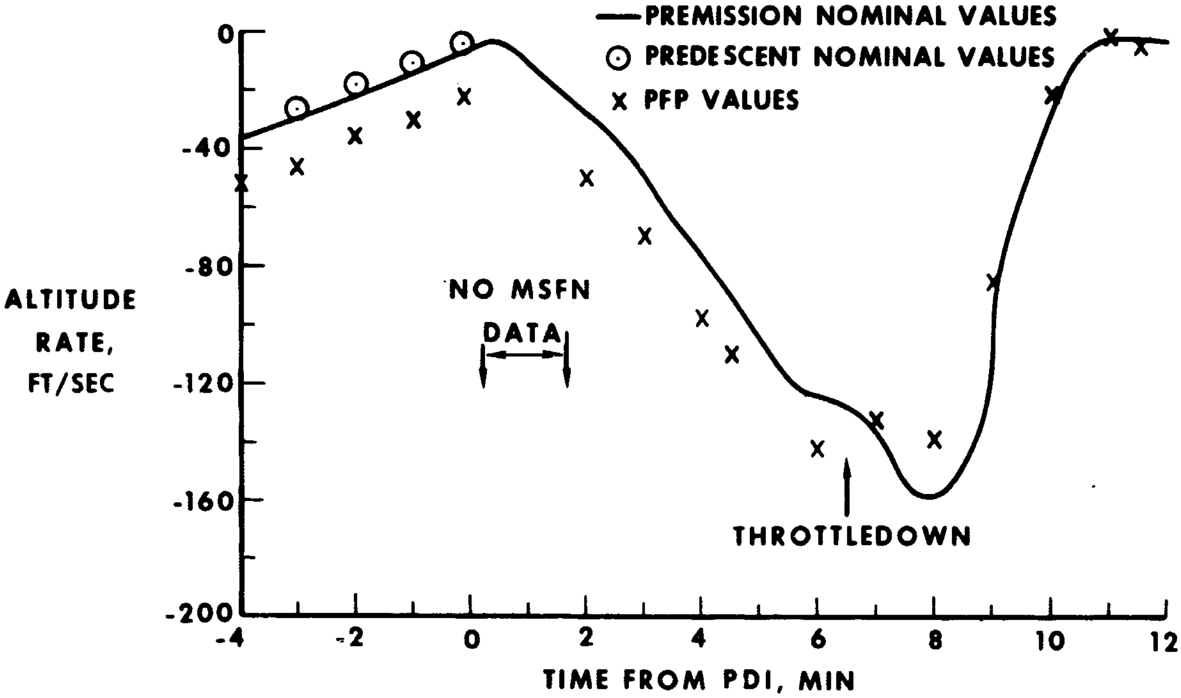

Figure 9. - Comparison of Apollo 11 premission and predescent nominal values with PFP values of LM altitude rate.

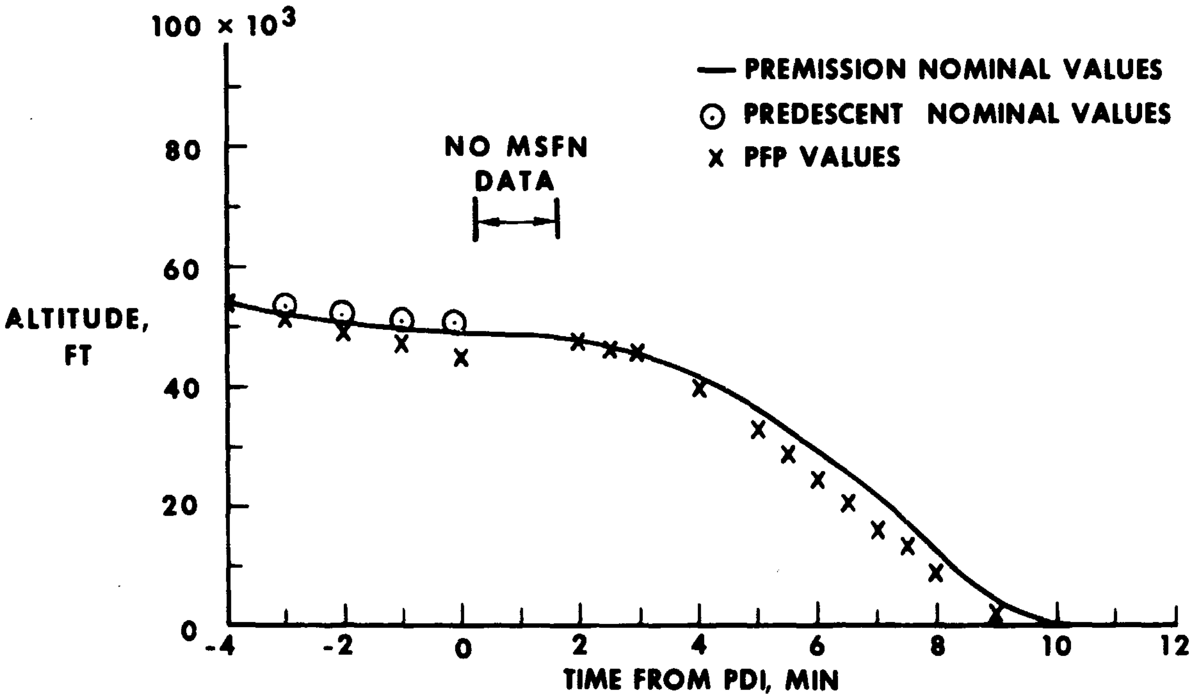

Figure 10. - Comparison of Apollo 11 premission and predescent nominal values with PFP values of LM altitude.

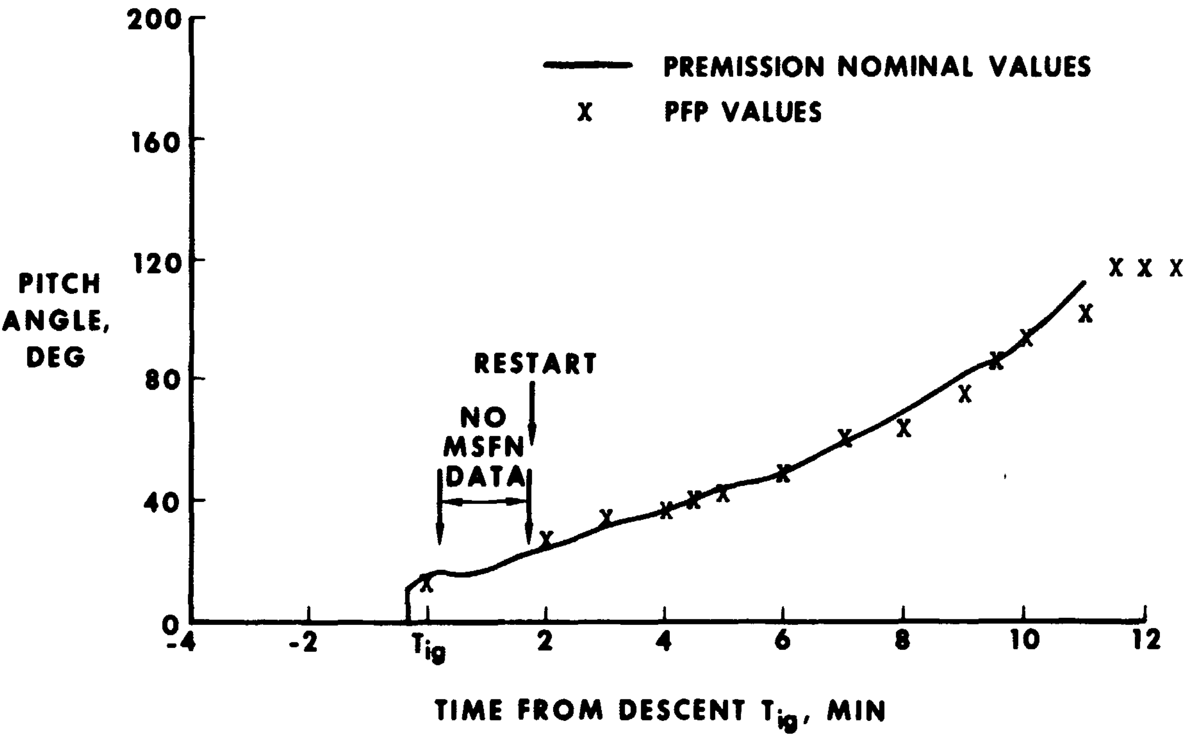

Figure 11. - Comparison of Apollo 11 premission nominal and PFP values of LM pitch angle.

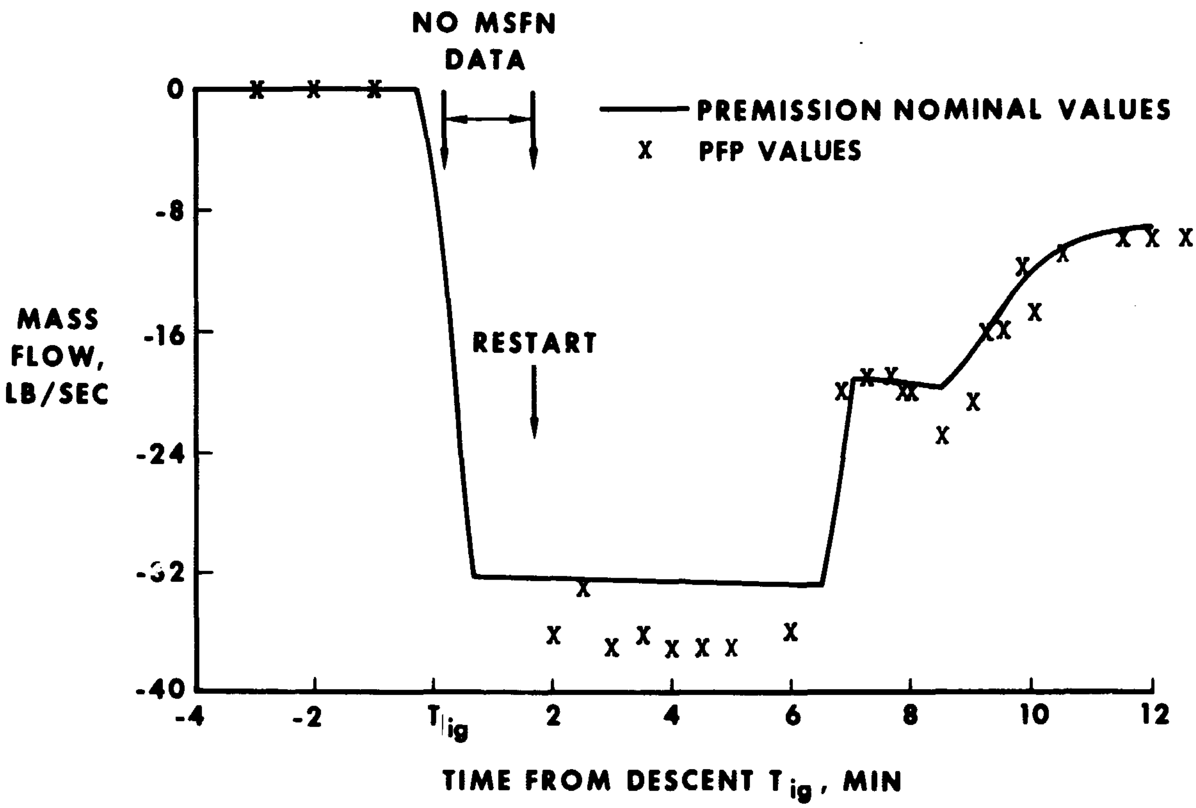

Figure 12. - Comparison of Apollo 11 premission nominal and PFP values of LM mass flow rate.

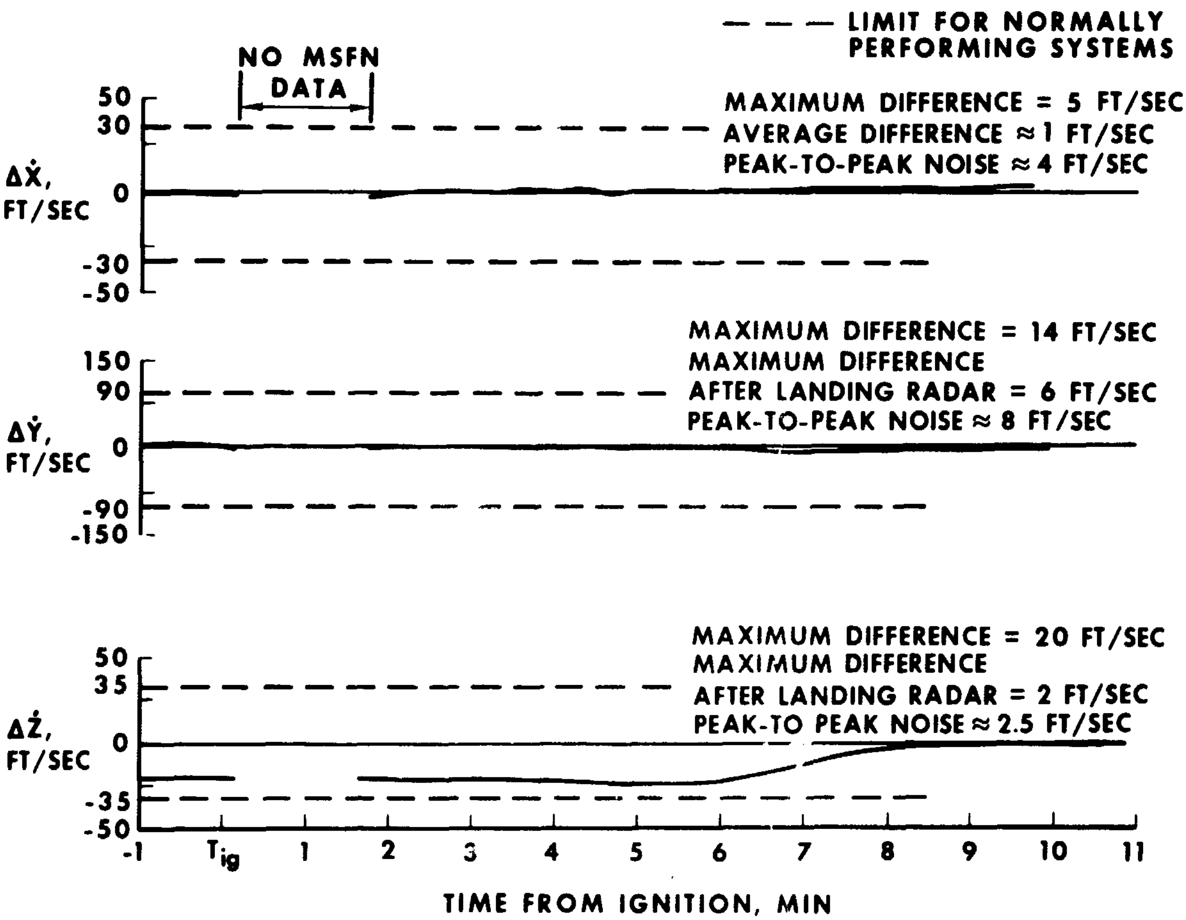

Figure 13. - Velocity comparison of PFP and PGNS during the Apollo 11 lunar descent.

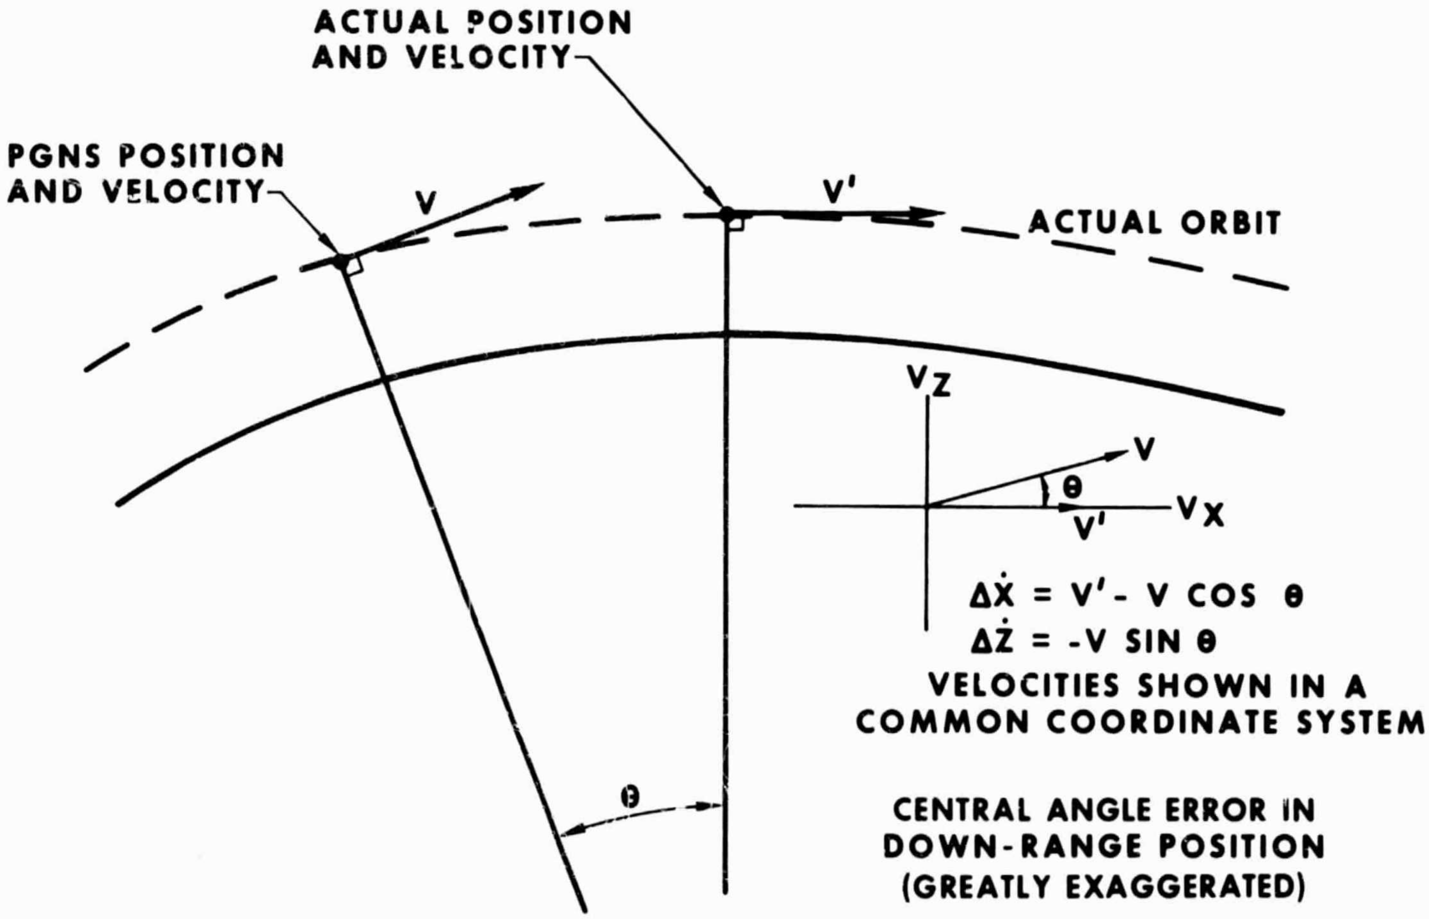

Figure 14. - Down-range position error effects on PFP and PGNS velocity comparison.

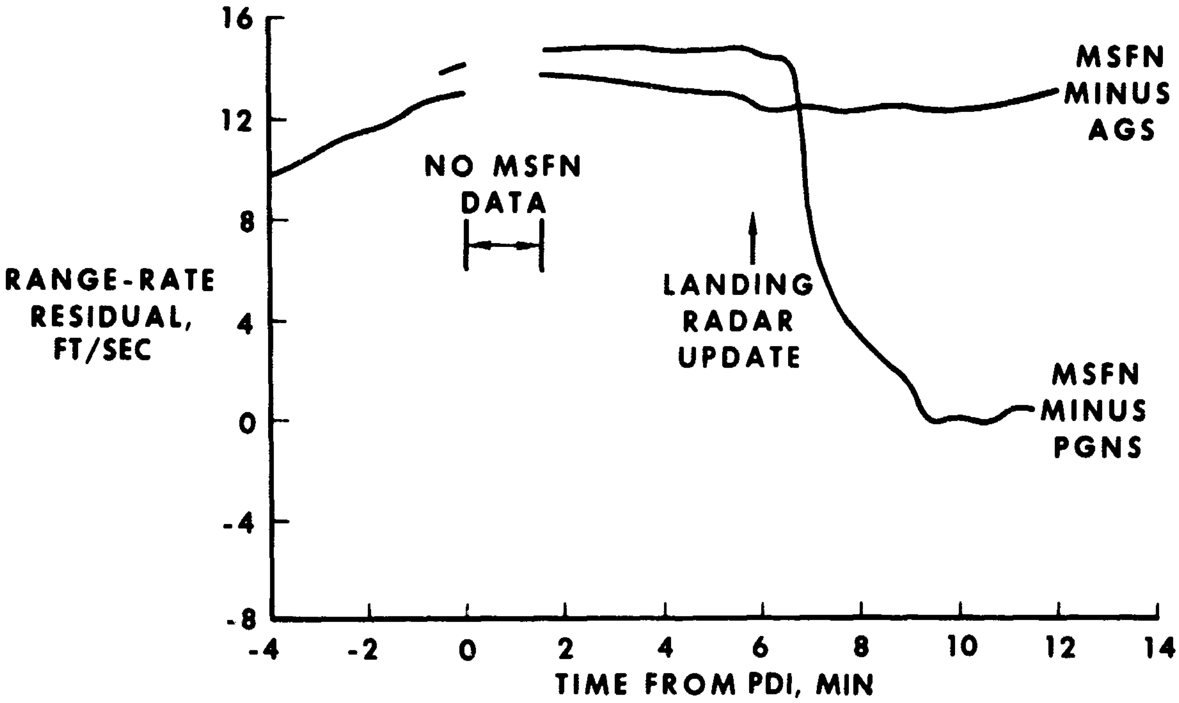

Figure 15. - Apollo 11 descent range-rate residuals.

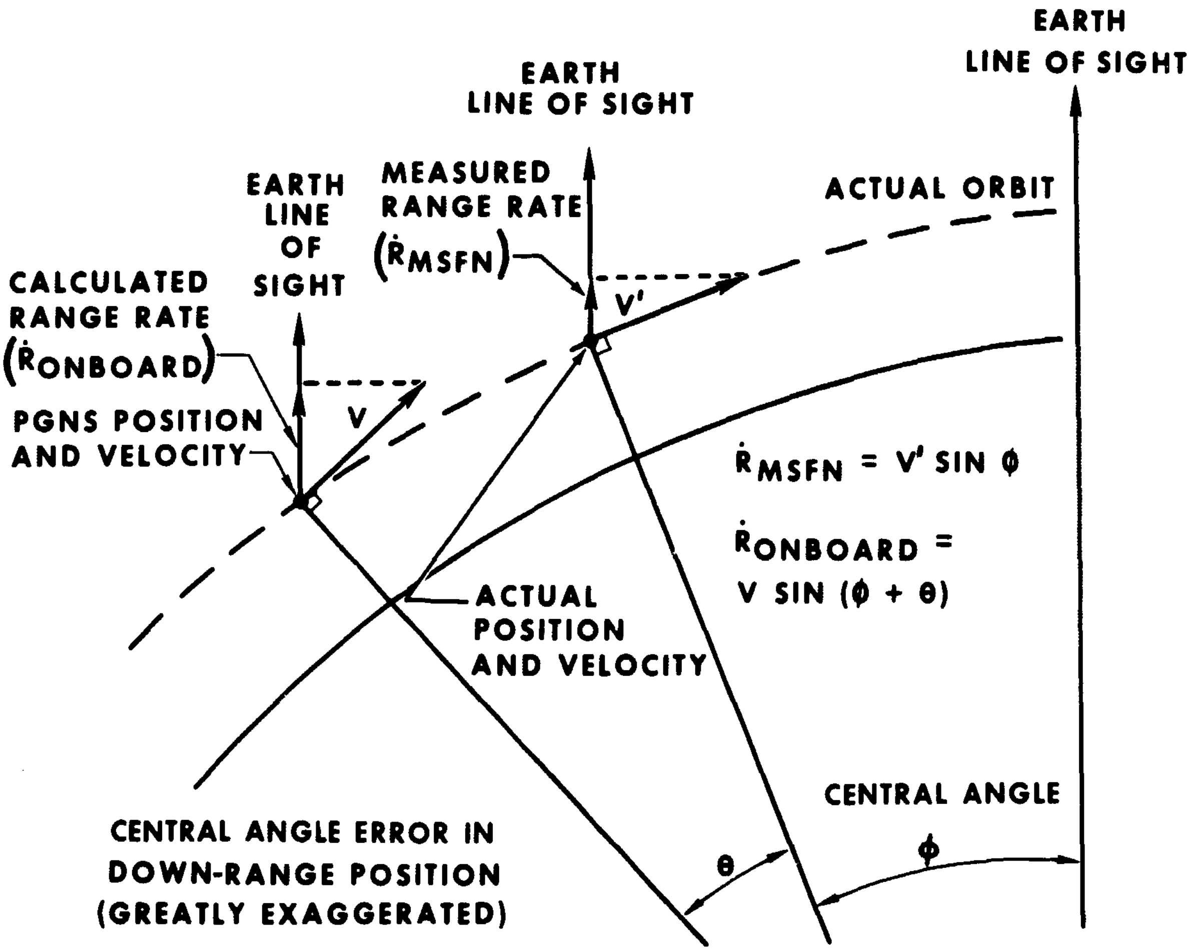

Figure 16. - Down-range position error effects on the MSFN range-rate processor.